

NO TIME FOR DEAD TIME: TIMING ANALYSIS OF BRIGHT BLACK HOLE BINARIES WITH *NuSTAR*

MATTEO BACHETTI^{1,2}, FIONA A. HARRISON³, RICK COOK³, JOHN TOMSICK⁴, CHRISTIAN SCHMID⁵, BRIAN W. GREFENSTETTE³, DIDIER BARRET^{1,2}, STEVEN E. BOGGS⁴, FINN E. CHRISTENSEN⁶, WILLIAM W. CRAIG^{4,7}, ANDREW C. FABIAN⁸, FELIX FÜRST³, POSHAK GANDHI⁹, CHARLES J. HAILEY¹⁰, ERIN KARA⁸, THOMAS J. MACCARONE¹², JON M. MILLER¹¹, KATJA POTTSCHMIDT¹³, DANIEL STERN¹⁴, PHIL UTTLEY¹⁵, DOMINIC J. WALTON³, JÖRN WILMS⁵, AND WILLIAM W. ZHANG¹⁶

Draft version September 12, 2014

ABSTRACT

Timing of high-count rate sources with the *NuSTAR* Small Explorer Mission requires specialized analysis techniques. *NuSTAR* was primarily designed for spectroscopic observations of sources with relatively low count-rates rather than for timing analysis of bright objects. The instrumental dead time per event is relatively long (~ 2.5 msec), and varies by a few percent event-to-event. The most obvious effect is a distortion of the white noise level in the power density spectrum (PDS) that cannot be modeled easily with the standard techniques due to the variable nature of the dead time. In this paper, we show that it is possible to exploit the presence of two completely independent focal planes and use the cross power density spectrum to obtain a good proxy of the white noise-subtracted PDS. Thereafter, one can use a Monte Carlo approach to estimate the remaining effects of dead time, namely a frequency-dependent modulation of the variance and a frequency-independent drop of the sensitivity to variability. In this way, most of the standard timing analysis can be performed, albeit with a sacrifice in signal to noise relative to what would be achieved using more standard techniques. We apply this technique to *NuSTAR* observations of the black hole binaries GX 339–4, Cyg X-1 and GRS 1915+105.

Subject headings: X-rays: stars — accretion, accretion disks — black hole physics

1. INTRODUCTION

Timing analysis, other than being an important diagnostic tool by itself (see Vaughan 2013, for a review), is particularly powerful for dissecting the inner accretion regions in black hole binaries and AGN when combined with spectral modeling. Good examples of this

combination in binaries are studies of the variability of a specific spectral component (e.g. the iron K_α line and the reflection component, Revnivtsev et al. 1999; Gilfanov et al. 2000), the study of the time lags between different energy bands and the comparison with the expected behavior of a given spectral component (e.g. Nowak & Vaughan 1996; Papadakis et al. 2001; Körding & Falcke 2004; Fabian et al. 2009; Gandhi et al. 2010; Uttley et al. 2011; Artigue et al. 2013), the study of the covariance of the signal at multiple energies and the variability spectrum (e.g. Uttley et al. 2011; Jin et al. 2013). A review of reverberation lag measurements in binaries and AGN can be found in Uttley et al. (2014).

The *Nuclear Spectroscopic Telescope ARray* (*NuSTAR*; Harrison et al. 2013) mission has deployed the first hard X-ray (above 10 keV) focusing satellite. It features two telescopes, focusing X-rays between 3 and 79 keV onto two identical focal planes (usually called focal plane modules A and B, or FPMA and FPMB). It has a field of view (FOV) of $12' \times 12'$ and an angular resolution of $18''$ FWHM; ($58''$ HPD). These features, together with good spectral resolution in the iron K band (~ 0.4 keV @ 6 keV), have motivated a large number of observations of X-ray binaries. A particularly interesting result was the accurate modeling of the reflection component in several black hole (BH) binaries (Miller et al. 2013; Tomsick et al. 2014) resulting in improved constraints on black hole spin in several systems. Such measurements have also been made by *NuSTAR* for active galactic nuclei (AGN; e.g. Risaliti et al. 2013; Marinucci et al. 2014; Walton et al. 2014). *NuSTAR*, *XMM* and *Suzaku* have recently measured reverberation lags in the AGN MCG-5-23-16 in both the iron line and Compton reflection hump together

matteo.bachetti@irap.omp.eu

¹ Université de Toulouse; UPS-OMP; IRAP; Toulouse, France

² CNRS; Institut de Recherche en Astrophysique et Planétologie; 9 Av. colonel Roche, BP 44346, F-31028 Toulouse cedex 4, France

³ Cahill Center for Astronomy and Astrophysics, Caltech, Pasadena, CA 91125

⁴ Space Sciences Laboratory, University of California, Berkeley, CA 94720, USA

⁵ Dr. Karl-Remeis-Sternwarte and ECAP, Sternwartstr. 7, 96 049 Bamberg, Germany

⁶ DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 327, DK-2800 Lyngby, Denmark

⁷ Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

⁸ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

⁹ Department of Physics, Durham University, South Road DH1 3LE, UK

¹⁰ Columbia Astrophysics Laboratory, Columbia University, New York, NY 10027, USA

¹¹ Department of Astronomy, University of Michigan, 500 Church Street, Ann Arbor, MI 48109-1042, USA

¹² Department of Physics, Texas Tech University, Lubbock, TX 79409, USA

¹³ CRESST, UMBC, and NASA GSFC, Code 661, Greenbelt, MD 20771, USA

¹⁴ Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

¹⁵ Anton Pannekoek Institute, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

¹⁶ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

for the first time (Zoghbi et al. 2014).

From the point of view of timing, *NuSTAR* has an advantage over other imaging satellites: the time resolution is $10\ \mu\text{s}$, and so one can in principle study variability over the whole range of interesting frequencies in accreting systems without switching to an observing mode with decreased spectral resolution. The satellite has in fact been used to perform timing analysis on selected targets with very good results: for example, it was used to study a number of accreting and rotation-powered pulsars, permitting the detection of variable cyclotron resonant scattering features (Fürst et al. 2013, 2014) and the discovery of strong hard time lags in the emission of another accreting pulsar (GS 0834-430; Miyasaka et al. 2013). *NuSTAR* was also used to detect the pulsations from a magnetar near the Galactic Center (Mori et al. 2013) and to carry out a detailed timing study of this serendipitous source (Kaspi et al. 2014). These pulsars are mostly slow rotators, with pulse periods above 1 s. Aperiodic timing has been measured in faint sources (e.g. Bachetti et al. 2013), or in bright sources but only over the low-frequency part of the power spectrum (e.g. Natalucci et al. 2014).

From the point of view of fast ($\nu \gg 1\ \text{Hz}$) timing studies of bright sources, however, *NuSTAR*'s default observing mode requires care to analyze properly. Each X-ray incident on a focal plane produces a trigger, and the energy signal is read out immediately. The read time is dependent on the number of triggered pixels, and so the dead time for each event is slightly variable (at the few percent level) with an average value of $\sim 2.5\ \text{ms}$. This dead time, together with additional dead time produced on a separate cadence by housekeeping operations and vetoed events, can be accurately measured and its value over a second is stored in the housekeeping files of the observations. In addition, the live time since the previous X-ray event is recorded in a column in the event lists (see Section 2 for details). For bright sources where the fraction of events vetoed by the anti coincidence system is small, light curves can be accurately produced with any time binning (see Madsen et al 2014, *submitted*, for application of this technique to the Crab pulsar). However, for timing analysis (e.g. the production of power density spectra), the dead time introduces spurious correlations between event arrival times that are reflected in a distortion of the power spectrum, making it difficult to subtract the Poisson noise level above $\sim 50\ \text{Hz}$. This effect is almost completely negligible for count rates up to $\sim 1\ \text{cts s}^{-1}$, but it is very pronounced in observations of Galactic X-ray binaries during outbursts, with typical incident count rates $\gtrsim 100\ \text{cts s}^{-1}$ over the full band.

In this Paper, we present a technique to avoid spurious noise introduced by dead time, in particular the distortion of the white noise level in the power spectrum. We show how a slight modification to the methods commonly used for quasi-periodic oscillation (QPO) detection and power spectral fitting enables good-quality results at all frequencies. We use these techniques to perform a basic timing analysis of several *NuSTAR* observations of BH binaries (BHBs), whose spectral analysis is presented in other papers (GRS 1915+105, Miller et al. 2013; Cyg X-1, Tomsick et al. 2014; GX 339-4, Fürst et al. *in prep.*). We chose these BHBs as ideal objects with which to demonstrate the timing techniques, since their timing

behavior has been extensively studied by other missions and instruments (see Section 6 for details).

The paper is organized as follows: in Section 2 we discuss details *NuSTAR*'s dead time, and how it is measured on-board the instrument; in Sections 3 and 4 we describe how to use the Cross Power Density Spectrum to obtain a proxy of the PDS; in Sections 5 and 6 we use these techniques to analyze the data from two black hole X-ray Binaries, Cyg X-1 and GX 339-4; and in Section 7 we provide conclusions. In the Appendix we present results from a timing analysis of GRS 1915+105, which was observed during the on-orbit commissioning period. This observation suffered from spurious features associated with the cadence of instrument housekeeping functions and so requires a unique analysis not relevant to *NuSTAR* science phase observations.

2. DEAD TIME IN *NuSTAR*.

Dead time in the *NuSTAR* instrument is produced when the focal plane module electronics are busy processing an event, when the shield veto prevents the focal plane from triggering (or interrupts an event that is being processed), and for some regular instrument housekeeping functions (the housekeeping produces less than 1 ms of dead time per second). *NuSTAR* has two independent telescopes, and the associated focal plane modules have independent processors, so that the dead time in the two modules is uncorrelated. All of the events processed by the instrument are found in the unfiltered event file. The cleaned event file is a subset of the events in the unfiltered event file where GTIs and additional screening have been applied to remove non-science quality events¹⁷.

For the production of light curves and flux measurements the *NuSTAR* instrument has two operating modes, one for faint ($< 50\ \text{mCrab}$) sources, and one that allows highly accurate flux measurements even at very high count rates and on arbitrarily short timescales. The fraction per second of dead time from events that are not processed, like shield vetoes, as well as that due to housekeeping functions is saved in the housekeeping files on a one-second cadence. For light curves binned in 1 s intervals (or a integer multiple thereof) multiplying by the live time stored in the housekeeping files correctly accounts for dead time. In addition, the live time since the previous X-ray event is stored in the event lists in the "PRIOR" column. In the absence of events vetoed by the active anti-coincidence shield, this column would accurately reflect the true live time. In the operating mode used for most observations taken since in-orbit checkout, the events vetoed by the anti-coincidence shield are not telemetered to the ground to minimize data volume. For this mode adding the PRIOR column does not yield the proper live time since vetoed events are not included. However, for source count rates significantly greater than the veto rate the error is negligible (for the Crab it is only $\sim 0.3\%$; see Madsen et al. 2014, *submitted*). For sources of intermediate brightness, where the veto rate cannot be ignored, a fully tested and calibrated instrument mode exists that includes vetoed events in the data stream so that fluxes and light curves can be produced with arbi-

¹⁷ See the NuSTARDAS user's guide available from the HEASARC (<http://heasarc.gsfc.nasa.gov/docs/nustar/analysis/>) for more information.

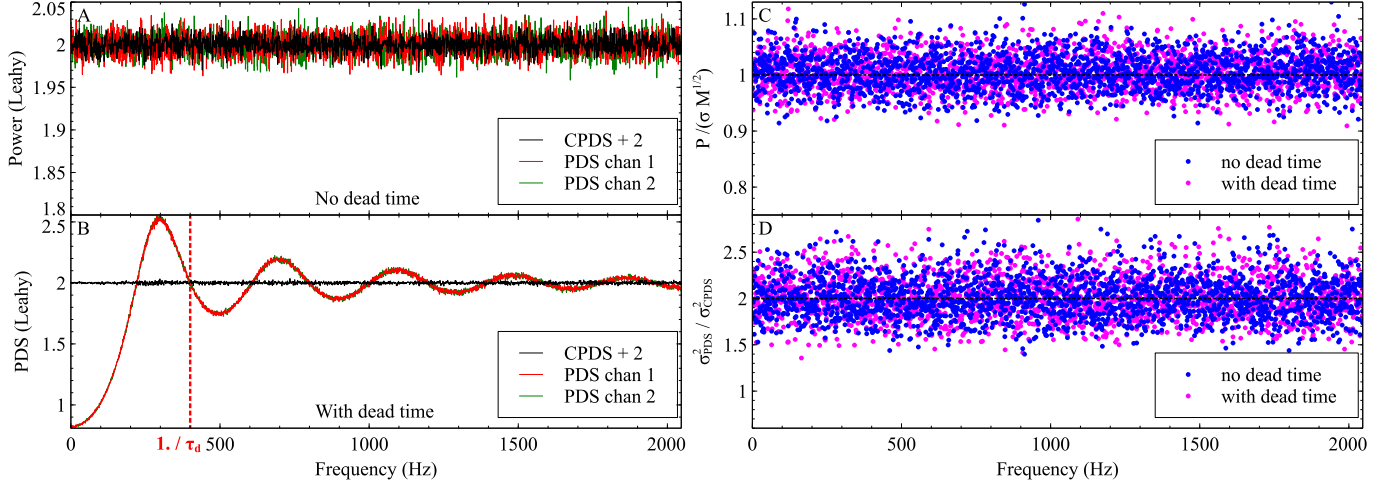


Figure 1. (Left) The cospectrum and the PDS are compared in the case of pure Poisson noise, without (A) and with (B) dead time. The simulated incident count rate was 225 ctss^{-1} . The cospectrum mean is always 0. In these plots, it has been increased by 2 for display purposes. The frequency $1/\tau_d$ is indicated. (Right) The usual relation between the PDS and its standard deviation ($\sigma = P/\sqrt{M}$, where M is the number of averaged PDSs) holds with and without dead time (C). Also, the variance of the cospectrum is half the variance of the PDS, in both cases (D).

trary accuracy.

For timing analyses that involve power spectral analysis, dead time produces systematic effects even if it is perfectly measured. Dead time is classified as either paralyzable or non-paralyzable depending on whether a photon hitting the detector during dead time produces new dead time or not. The *NuSTAR* readout architecture produces non-paralyzable dead time. Its main effect is on recorded count rates: if we have an incident rate of photons r_{in} , the detected count rate will be approximately

$$r_{\text{det}} = r_{\text{in}} / (1 + \tau_d r_{\text{in}}) \quad (1)$$

(see, e.g., van der Klis 1989, and references therein), where τ_d is the dead time produced by each event, assuming that it is constant. But dead time also alters the sensitivity to variable signals, acting as a frequency filter. Power density spectra (PDS), in particular, are deformed to a “wavy” shape that depends on the magnitude of dead time and on count rates (see Figure 1). Power at frequencies slightly above $1/\tau_d$ is *quenched*, as there is a lack of events whose separation is less than τ_d , while there is a relatively higher rate above τ_d , and therefore the power at frequencies just below $1/\tau_d$ is slightly *amplified*. These “waves” have nodes at $1/\tau_d$ and multiples thereof, where the power (in Leahy et al. (1983) normalization) is equal to 2, the value that it would have without dead time, and maxima and minima in between are given by the relative contribution of the quenching and amplification. For frequencies $\nu \ll 1/\tau_d$, the main effect is a general deficiency of events, and the power has a decreasing level that approaches $\approx 2(1 - r_{\text{in}}\tau_d)^2$ (Weisskopf 1985). Assuming that τ_d is the same for each event and that only source events contribute to either non-paralyzable or paralyzable dead time, this distortion can be modeled precisely (see Vikhlinin et al. 1994; Zhang et al. 1995, for an exhaustive treatment). Also, some statistical properties of the PDS hold in dead time-affected data. For example, the standard deviation associated with the bin i of the PDS is always equal to P_i/\sqrt{M} , where P_i is the power in the bin i and M is the number of averaged PDSs (see Figure 1).

In *NuSTAR* τ_d is not strictly fixed at the same value for all events, but varies by a few percent depending on the number of pixels that are triggered. For this reason, the models available in literature do not correctly describe the dead time effects for this satellite: the “wavy” behavior of the PDS shifts slightly, and to fully account for this effect and produce a white-noise subtracted PDS, a very precise modeling of the dead time would be required. Since at high count rates the “waves” can be very prominent, any real variability feature such a QPO can easily be “hidden” and difficult to detect.

As an additional complication, the models described above assume that dead time is produced completely by the recorded signal. In *NuSTAR*, additional dead time comes from events outside the source extraction region, from vetoed events, and from all events discarded for other reasons during the cleaning process in the pipeline (the step from unfiltered to cleaned event files). In the following, we neglect the effect of vetoed events, since their dead time ($\sim 20 \mu\text{s}$) has characteristic frequency $1/\tau_d \sim 50 \text{ kHz}$, much higher than science events, and their total contribution to dead time is small. We instead present a method that permits construction of a proxy of a white-noise subtracted PDS, regardless of the count rate and the ratio between source and background (or spurious) events.

3. THE CROSS POWER DENSITY SPECTRUM AS AN IMPROVED POWER SPECTRUM

NuSTAR has two completely independent focal plane modules (each containing four detectors) that are read out by separate microprocessors. It is therefore possible in principle to obtain the same information given by a PDS through the so-called Cross Power Density Spectrum (CPDS; for more details see Bendat & Piersol 2011): instead of considering the PDSs in the two individual focal planes:

$$P_i(\nu) = \mathcal{F}_i^*(\nu) \mathcal{F}_i(\nu) \quad (i = A, B) \quad (2)$$

where \mathcal{F}_i indicates the Fourier transform of the lightcurve detected by the focal plane i and ν is the frequency, one

multiplies the complex conjugate of one Fourier transform with the other Fourier transform:

$$C(\nu) = \mathcal{F}_A^*(\nu) \mathcal{F}_B(\nu). \quad (3)$$

The CPDS is often used in other contexts to obtain information on the correlation between the signal in two energy bands. It is a complex quantity: its real part is also called the *cospectrum* and gives a measure of the signal that is *in phase* between the two channels; its imaginary part, or *quadrature spectrum*, gives instead a measure of the *off-phase* signal. Therefore, in principle, it should be possible to eliminate all variability that is not related between the two light curves, including the effects of dead time, by just considering the cospectrum (the real part of the CPDS). In Figure 1 we show the statistical properties of the cospectrum in the case of pure Poisson noise. In both the dead time affected and in the zero-dead time cases the cospectrum mean value is 0 (in Figure 1, it has been shifted to 2 for graphical reasons). This is a big advantage, as this is independent of whether the dead time is constant or not (since the distribution of dead time is also independent between the two detectors), and therefore it is not necessary to conduct complicated studies of the dead time distribution in order to obtain a white noise-subtracted cospectrum, as opposed to a PDS where this procedure would be needed.

Another important point is that the standard deviation σ_{CPDS} of the cospectrum is linked to the standard deviations (σ_{PDS}) of the PDSs of the two light curves. In Figure 1, we show that the ratio between σ_{CPDS} and the geometric average of the standard deviations of the two PDSs is close to $\sqrt{2}$ regardless of the frequency. This is very convenient, as it allows us to assign to the cospectrum bin c_i an uncertainty of $\bar{P}_i/\sqrt{2M}$, where \bar{P}_i is the geometric average of the two PDSs and M is the number of averaged PDSs. The argument (or angle) of the CPDS is called *phase lag*; if we divide it by $2\pi\nu$ we obtain the *time lag* t_{lag} , that is a measure of the time shift between the two channels. In our case, since we are using two lightcurves in the same energy bands, we expect any time lag to be zero. But time lags between different energy bands have been used as indicators of how a signal produced in an emitting region is reprocessed in other regions: for example, if the disk emission is Comptonized in a corona, or when the signal from a region is reflected or propagated in another region, this time delay between the different emissions can be detected through time lags (see, e.g., Bendat & Piersol 2011; Nowak et al. 1999a; Uttley et al. 2014 and references therein). We briefly discuss time lags between different energy bands in the following sections, but for the most part our analysis will be performed on two channels in the same energy band, and the only important quantity for our treatment will be the cospectrum.

4. SIMULATIONS

In order to provide a fully consistent treatment of dead time effects in the PDS and in the cospectrum, we ran a large number of simulations. In each simulation we produced two event series containing variability, one for each FPM, and analyzed the data with cospectrum and PDS before and after applying a dead time filter. By doing so, we studied in detail the properties of the cospectrum




Figure 2. QPOs of equal Q factor (20) and r.m.s. amplitude (8%) at different frequencies. The units in this plot, and in all following power spectra or cospectra if not stated otherwise, are $\text{Power} \times \text{Frequency}$. Grey points show the standard, dead time-affected PDS. The white noise level to subtract in the dead time-affected datasets was calculated between 10 and 20 Hz (minimum between two QPOs), while in the dead time-free case we subtracted the theoretical level (2 in Leahy normalization).

and compared them to those of the PDS. In the following paragraphs we will explain the procedure in more detail, and demonstrate that the cospectrum can be considered a very good proxy of a white noise-subtracted PDS, albeit with some corrections to account for in the measured r.m.s..

4.1. Procedure

Light curve generation— We used the procedure by Davies & Harte (1987), introduced to astronomers by Timmer & Koenig (1995), to simulate light curves between two times t_0 and t_1 , from a number of model PDS shapes containing QPOs. The sampling frequency of the light curves was at least four times higher than the maximum frequency of the variability components included in the simulation. We normalized the light curves in order to have the desired mean count rate and total r.m.s. variability (7-10%). In order to be able to later calculate PDSs with a given maximum timescale T (see “Calculation of the PDSs and the cospectrum” below), we simulated light curves at least 10 times longer, following the prescriptions in Timmer & Koenig (1995) to avoid aliasing.

Event list generation— From every light curve, we generated two event lists, corresponding to the signal from the two focal planes. Each event list was produced as follows: first of all, we calculated the number N_{save} of event times to be generated as a random sample from a Poisson distribution centered on the number of total photons expected (summing up all expected lightcurve counts); then, we generated N_{save} events by using a *Monte Carlo acceptance-rejection method*. This is a classical Monte Carlo technique; a more general treatment can be found in most textbooks on Monte Carlo methods (e.g. Gentle 2003). In our case, we used the following procedure: (1) for every event we simulated an event time t_e uniformly distributed between t_0 and t_1 , and an associated random amplitude (“probability”) value A_e between 0 and the maximum of the light curve; (2) we rejected

all t_e whose associated A_e was higher than the light curve at t_e ; to avoid possible spurious effects given by a stepwise model light curve, we used a *cubic spline* interpolation to approximate the light curve between bins; (3) we sorted the event list by t_e . To simulate the effects of background (spurious events filtered out by the pipeline, events recorded outside the source regions, etc.), we also produced two background series, at constant average flux, one for each simulated source light curve.

Dead time filtering— For each event list, we created a corresponding dead time-affected event list by applying a simple dead time filter: for each event, we eliminated all events in the 2.5 ms after it. Source and background events contributed equally to dead time. We used different versions of the dead time filter by varying slightly the dead time between events (of ~ 0.1 ms). But the pernicious effect of variable dead time is mainly on white noise subtraction. In our case (Figure 1), the cospectrum allows us to overcome this problem as its white noise is 0, and we verified that the other effects are not significantly different in the constant and variable dead time cases. In the following, we will treat the case with constant dead time.

Calculation of the PDSs and the cospectrum— We divided each pair of event lists in segments of length T , calculated the PDS in each of segments and the CPDS from each pair of them. We then averaged the PDSs and CPDSs from all segments. The CPDS white noise level is already 0. For the PDSs, used only in the ideal zero-dead time case, we subtracted the theoretical Poisson level (2 in Leahy normalization). We then rebinned the PDSs and CPDSs, either by a fixed rebin factor or by averaging a larger number of bins at high frequencies, following approximately a geometric progression. We finally multiplied the Leahy-normalized PDSs by $(B+S)/S^2$ (where B is the background count rate and S the mean source count rate) in order to obtain the squared r.m.s. normalization often used in literature (Belloni & Hasinger 1990; Miyamoto et al. 1991). The CPDS was instead multiplied by the factor $(\bar{B} + \bar{S})/\bar{S}^2$, where bars indicate the *geometric* averages of the count rates in each of the two event lists used to calculate it¹⁸.

Finally, we calculated the cospectrum by taking the real part of the CPDS. As described above, we assigned to each final bin of the cospectrum c_i an uncertainty calculated from the geometrical average of the PDSs in the two channels, divided by $\sqrt{2MW}$, where M is the number of averaged spectra and W is the number of subsequent bins averaged to obtain c_i .

Fitting procedure— Cospectra do not need Poisson noise subtraction; for PDSs, we fitted a constant to the interval outside of the frequency range containing QPOs. Since in the following paragraphs, in all our examples, we will be showing power spectra obtained by averaging more than 50 PDSs, we are in the Gaussian regime and a fit

¹⁸ We chose this normalization by analogy: if the PDS is multiplied by $(B+S)/S^2$, it's as if each Fourier transform contributing to it were multiplied by $((B+S)/S^2)^{1/2}$. Since the contribution to the CPDS is by two separate Fourier transforms, we multiplied by $((B_A + S_A)/S_A^2)^{1/2}((B_B + S_B)/S_B^2)^{1/2}$, that is equivalent to the geometric averages described.




Figure 3. (Top) Variation of the r.m.s. of a QPO at different peak frequencies, measured with the various techniques and with and without dead time. Each point represents a simulated QPO with r.m.s. = 10% and FWHM = 2 Hz. A total of 281 simulations were used for this plot. (Bottom) Significance measured with each method. The total PDS has about twice the significance of the single-module PDS in the no-dead time case, as expected, due to double number of photons. The CPDS in the no-dead time case is a factor $\sim \sqrt{2}$ higher than the single PDS, and lower by the same amount than the total PDS. The dead time-affected CPDS, instead, has a much lower level due to the lack of photons. The decrease of significance does not depend on the frequency of the QPO, but only on count rate (see Figure 4).

with standard χ^2 minimization routines is appropriate to the precision we are interested in (van der Klis 1989; Barret & Vaughan 2012).

Then, we fitted the QPOs with a Lorentzian profile in XSPEC¹⁹ (Arnaud 1996). Errors were calculated through a Monte Carlo Markov Chain, as those intervals where the χ^2 of the fit with no frozen parameters increased by 1. According to Lewin et al. (1988), the significance of the detection of QPOs is expected to be

$$n_\sigma \simeq \frac{1}{2} \frac{S^2}{B+S} r^2 \left(\frac{T}{\Delta\nu} \right)^{1/2} \quad (4)$$

where r is the r.m.s. and $\Delta\nu$ is the equivalent width of the feature (for a Lorentzian, $\Delta\nu = \pi/2 \times \text{FWHM}$). The exact proportionality factor depends on the definition of significance. In our case, the significance of QPOs was defined as the ratio between the amplitude of the Lorentzian and its error. The significance calculated in this way gives a value ~ 2 times lower than obtained by using the excess power à la Lewin et al. (1988) (a factor 2 is expected due to the fact that the Gaussian errors are calculated over \mathbf{R} , and the excess power over \mathbf{R}^+ ; see Boutelier 2009; Boutelier et al. 2009), but the trend in Equation 4 holds provided that the variability is dominated by Poisson noise.

¹⁹ We used both the standard interface of the program, and its Python bindings, PyXSPEC.




Figure 4. Similar to Figure 3, but with the centroid frequency of the QPO fixed at 30 Hz and letting the count rate vary between 10 and 1000 cts s^{-1} . A total of 118 simulations were used in this plot. The line shows Equation 5. It is not a fit, and describes remarkably well the data. In the bottom panel, we plot the detection significance for all the cases. The lines, again, are not fitted, they just show the theoretical prediction from Equation 4. All dead time-free cases are in good agreement with a linear increase with count rate below $\sim 500 \text{ cts s}^{-1}$, above which some curvature appears due to the departure from the quasi-Poissonian regime. The dead time-affected case is in good agreement with Equation 4 if, instead of the incident count rate (dashed line), one uses the *observed* count rate (solid line; see Equation 5).




Figure 5. The dependence of the r.m.s. on the relative contribution of the source in a given energy range to the total count rate. The 132 simulations that compose this plot show how the r.m.s. drop is stronger if the source signal dominates the total signal, since the source signal contributes more to the total dead time. The solid black line shows Equation 6. It is not a fit.

4.2. Simulation results

First look— The simulation in Figure 1 shows the comparison between the PDS and the cospectrum with and without dead time for pure Poisson noise. From these simulations it is clear that the white noise-subtracted PDS and the cospectrum are equivalent in the case with

no dead time. It is immediately evident that the most problematic effect of dead time, the modulation of the white noise level, disappears in the cospectrum. In the following paragraphs, we investigate the frequency and count rate dependence of these quantities in more detail.

Frequency dependence— The general statistical properties of the PDS and the cospectrum are also very similar, both in the dead time-affected and in the zero-dead time case. Figure 1, panel D, shows that the variance of the cospectrum and the PDS maintains a constant ratio equal to 2 in both the clean and the dead time-affected datasets. This makes it easy to calculate the variance of the cospectrum values for the subsequent analysis, by simply using the known properties of the PDS where the variance is just equal to the square of the power (in Leahy normalization).

Contrary to what might be imagined, it is possible to detect variability even at frequencies that are affected the most by dead time, i.e. those above $1/\tau_d$. Figure 2 shows that QPOs at all frequencies are detectable, albeit with some modulation of the observed r.m.s.. To measure this change of r.m.s., we simulated ~ 500 lightcurves using the method above, each containing a single QPO with frequencies equally distributed between 5 and 1000 Hz, r.m.s. = 10% and FWHM = 2 Hz. As explained above, from every light curve we obtained two event lists in order to simulate the signals from the two detectors. We produced for every pair of event lists the cospectrum, the two PDSs, and a total PDS including the counts from both detectors, both in the zero-dead time and the dead time-affected cases. We then fitted the resulting spectra with a Lorentzian model in XSPEC.

Figure 3 shows the change of the r.m.s. measured with the PDS and the cospectrum, with and without dead time. The measured r.m.s. in the zero-dead time PDSs agrees with the zero-dead time cospectrum, whereas the dead time-affected cospectrum yields a frequency-dependent deviation from the true r.m.s., with deviation following the same trend as the variance (see also Figure 1).

The detection significance does *not* depend on the frequency in any case, with or without dead time. The decrease of the significance is instead driven by the observed count rate, as we discuss shortly. In the no-dead time case, the significance of the single PDSs is about half that of the total PDS because the significance is directly proportional to the intensity of the signal (Equation 4), and in the total PDS one uses twice the number of photons. The zero-dead time cospectrum yields instead a significance $\sim \sqrt{2}$ lower than the total PDS, and higher by the same amount than the single-module PDS. This is just an effect of the factor $\sqrt{2}$ between the standard deviations of the cospectrum and the single-module PDS. From Figure 3, it is clear that it is advantageous to use the total PDS for low count rates where dead time is negligible, and the cospectrum otherwise, but with the formulas and the simulations shown above to account for the frequency-dependent distortion of the r.m.s. amplitude.

In summary, the important point that Figure 3 makes is that QPOs are still detectable at any frequency, even those heavily affected by dead time, albeit with a change of the measured r.m.s. that must be taken into account.

Count rate dependence— We now investigate how the measured r.m.s. is influenced by count rate. Figure 4 and 5 show the variation with count rate in the detected r.m.s. of a QPO at 30 Hz, FWHM=2 Hz and r.m.s. =7.5%, in two cases: increasing total count rate, and fixed total (source + background) count rate with variable source count rate. For the first case, since $30 \text{ Hz} \ll 1/\tau_d$ and the higher-order corrections are not needed, we use van der Klis (1989; equations 3.8 and 4.8) to obtain

$$\frac{r.m.s._{det}}{r.m.s._{in}} \approx \frac{1}{1 + \tau_d r_{in}} = \frac{r_{det}}{r_{in}}. \quad (5)$$

This relation is plotted with a dashed line in the top panel of Figure 4 and it is in remarkably good agreement with the simulated data.

In general, one would expect the significance of detection in the PDS to be proportional to the incident count rate and to the square of the r.m.s. (Equation 4). This condition holds if the QPO can be considered a small disturbance in an otherwise Poissonian process, or $r.m.s. \ll \sqrt{\nu/r_{in}}$ (Lewin et al. 1988). In the bottom panel of Figure 4, we fit Equation 4 below 600 cts s^{-1} , with a multiplicative constant due to the slightly different definition of $1\text{-}\sigma$ error that we use ($\Delta\chi^2 = 1$ instead of the Leahy et al. (1983) definition). The best-fit multiplicative constant, $\sim 1/2.2$, turns out to be consistent with the factor 2 expected from the fact that we are using Gaussian fitting instead of excess power (see Section 4.1). The departure from the linear condition above $\sim 600 \text{ cts}$ is evident. Indeed, it is expected that at count rates above $0.1 \nu/r.m.s.$, the significance starts departing from the linear trend. The total PDS is visibly more affected because its count rate is double than that of the single-module PDS. The significance of detection with the cospectrum is $\sim \sqrt{2}$ lower than that of the total PDS, while it is $\sim \sqrt{2}$ higher than that of the PDS from a single module. The dead time-affected cospectrum, instead, has a large deviation from the linear trend. This is just an effect of the diminishing count rate due to dead time. In fact, what is plotted is the *incident* count rate. If one converts it to the *detected* count rate, the linear relation between count rate and significance still holds (solid line).

The second case (Figure 5) clearly shows a linear decrease of the measured r.m.s. as the source gains counts with respect to the background. Again, by using van der Klis (1989; equations 3.8 and 4.8), but this time putting the *total* count rate ($r_{in} + r_{back,in}$, where $r_{back,in}$ is the non-source count rate) in the relation between incident and observed count rates, one obtains

$$\frac{r.m.s._{det}}{r.m.s._{in}} \approx \left[1 - \frac{\tau_d r_{in}}{1 + \tau_d (r_{in} + r_{back,in})} \right] \quad (6)$$

This means that the measure of r.m.s. we obtain in our data will be generally affected more if the source signal dominates the background, as is the case in most *NuSTAR* observations of bright sources. In the examples that we present in the next Sections, we make use of Monte Carlo simulations similar to the ones above to estimate the change of r.m.s. at the count rate of the sources we observe.

Table 1
Summary of the observations used in this work

Source	ObsId	Date y-m-d	On time ks	Livetime ks
GRS 1915+105	10002004001	12-07-03	27.0	15.7
Cyg X-1	30001011002	12-10-31	18.4	10.8
Cyg X-1	30001011003	12-10-31	10.3	5.1
GX 339-4	80001013002	13-08-12	38.8	35.1
GX 339-4	80001013004	13-08-16	23.4	20.8
GX 339-4	80001013006	13-08-24	26.8	23.2
GX 339-4	80001013008	13-09-03	11.1	9.4

5. OBSERVATIONS AND DATA REDUCTION

We now demonstrate the methods outlined above by analyzing the power density spectra for a number of bright Galactic black hole binaries.

A summary of the observations used in this work is listed in Table 1. We preprocessed all *NuSTAR* observations with the NUSTARDAS pipeline included in HEASOFT 6.15. We produced cleaned event files and calculated good time intervals (GTIs). We processed both the unfiltered and the clean event files (files tagged with “_uf” and “_cl” respectively in the “event_cl” directory of *NuSTAR* data). By analyzing the unfiltered and cleaned event files, it is possible to estimate the ratio of “good” events to the total of dead time-producing events.

We calculated the ratio r_{good} of “good” events over intervals of 100s. We found that in all observations r_{good} was almost constantly $0.9 \lesssim r_{good} \lesssim 0.94$, with some drops due to increased solar activity. We selected the intervals in which this ratio didn’t go below a certain threshold (ranging from 0.9 to 0.92 in different files), in order to eliminate possible spurious effects on the observed variability. In the subsequent analysis, we will always consider an additional 10% “background” event rate due to unrecorded and spurious events.

We selected events from a region around each source depending on how broad the point-spread distribution was (see the next subsections for details for each source).

We used the **barycorr** FTOOL to calculate event times at the solar system barycenter using the DE200 ephemeris and source positions from Simbad²⁰.

Additional variability is introduced into the cross-power spectrum in some cases due to mis-matched GTI between FPMA and FPMB. This mis-match occurs often when the satellite traverses the South Atlantic Anomaly (SAA) because the automatic algorithm that turns off the focal plane modules to conserve telemetry kicks in at different thresholds for FPMA and FPMB. The parameters that rule this behavior are finely tuned in order to maximize the number of photons available for spectral analysis while avoiding most of the contamination (Harrison et al. 2013), but for timing analysis this produces unwanted effects that have to be accounted for. A signature of this variability is a strong time lag between the two detectors in the same energy band (therefore, unrelated to the physics of the source), that disappears when discarding an additional interval at the GTI borders or by using a more aggressive filtering during the pipeline process. We decided to use a “safe” additional cut of 300s at the borders of all GTIs during timing analysis, and verified that it solved the problem in all affected

²⁰ <http://simbad.u-strasbg.fr/>

ObsIds.

6. DATA ANALYSIS

From each barycentered event list (Section 5) we obtained light curves sampled at $1/2048\text{s}$ ($\sim 488\mu\text{s}$). We extracted separate light curves for each detector, and calculated the common Good Time Intervals (GTIs).

Each light curve was then divided in segments (the length was chosen in order to optimize the use of GTIs) with the start and stop times locked between the two detectors, and a Fourier transform was calculated from each of them. We then calculated the CPDSs from each segment by multiplying the Fourier transforms as explained in Section 2, and obtained a single cospectrum from the real part of the average of all CPDSs.

When needed, we plot the error on the best-fit model as a hatched region around the model line. To calculate these errors, we ran a Monte Carlo Markov Chain on XSPEC to simulate 10000 models around the best fit. We chose all the models having $\Delta\chi^2 < 2.706$, and we calculated at each energy the maximum deviation from the best fit coming from any of these models.

6.1. *Cyg X-1*

Cyg X-1 is a bright persistent BHB, discovered during one of the earliest rocket flights at the dawn of X-ray astronomy (Bowyer *et al.* 1965). As such, it is among the best studied X-ray sources in the Galaxy. Its spectral evolution is typical for a BH, with the succession of low/hard, high/soft and intermediate states with correspondingly different timing patterns (see Dunn *et al.* 2010, for a broad view on BHBs). Variability in all spectral states has been observed. It shows two dominant Lorentzian components in the low/hard state. When the source softens their characteristic frequencies increase and an additional cutoff power law-like component appears. In the soft state the latter dominates the power spectrum (Nowak & Vaughan 1996; Nowak *et al.* 1999a; Revnivtsev *et al.* 2000; Churazov *et al.* 2001; Axelsson *et al.* 2005; Böck *et al.* 2011, and references therein). Time lags have been observed in all spectral states (Pottschmidt *et al.* 2000); the variability of single spectral components was also investigated, showing that the reflection component shows variability at lower frequencies than the bulk of the emission in the low/hard state (Revnivtsev *et al.* 1999) while the r.m.s. level and shape of the power spectrum is comparable at all energies in the soft state (Gilfanov *et al.* 2000). A very thorough investigation of the relation between variability in different energy ranges and spectral states can be found in Grinberg *et al.* (2014). For past studies of the high-energy variability of the source see Cabanac *et al.* (2011).

NuSTAR observed Cyg X-1 in its soft state (for a full spectral analysis description, see Tomsick *et al.* 2014). The spectral analysis shows a strong reflection component, that might produce detectable phase lags. Phase lags have been observed with *RXTE* in all spectral states (Miyamoto & Kitamoto 1989; Pottschmidt *et al.* 2000; Böck *et al.* 2011; Grinberg *et al.* 2014; see Nowak & Vaughan 1996; Nowak *et al.* 1999a for a rigorous treatment).

The PDS in this state is often well described by a cutoff power law-like shape almost “flat” between 10^{-2}

and 10 Hz (e.g. Cui *et al.* 1997; Churazov *et al.* 2001; Gleissner *et al.* 2004; Axelsson *et al.* 2005). This is the case in our observations, as shown in Figure 6. There is a clear change in r.m.s. variability between the two observations, and a power law-like red noise component appears at the lowest Fourier frequencies in only one observation, but for the rest the cospectrum is very well described by a cutoff power law.

As shown in Figure 4, a drop of r.m.s. is expected at higher count rates, and the second observation has indeed a higher count rate than the first. But in Figure 6 we show that the expected r.m.s. drop due to dead time is far lower than observed. This means that the rms change is largely due to an intrinsic lower variability of the source in the second observation.

More interesting is the cospectrum in different energy ranges, shown in Figure 6. There is a very strong change in the shape and normalization of the cospectrum with energy. In general, the r.m.s. increases with energy, and it is about an order of magnitude higher above 10 keV than below 5 keV . This is a clear indication that the soft component of the spectrum (coming from the disk) is varying less than the reprocessed component (reflected and Comptonized).

Again, in order to evaluate the r.m.s. change with energy, it is necessary to apply a correction for dead time. In fact, channels with more counts will have a larger drop of r.m.s. than channels with a few counts (compare to Figure 5).

Following the procedure in Section 4, we simulated the expected relative drop in r.m.s. in all bands. In Figure 6 we show the measured r.m.s. and the curves showing the expected drop due to count rate. These curves were calculated with the method depicted in Figure 4, assuming a total count rate ~ 1.1 times the source count rate, as the total background plus the rejected events account for about 10% of the contributions to dead time. As can be seen, the r.m.s. change between bands is mostly intrinsic, with a drop of only $\sim 10\%$ expected for the channels with the highest count rates. Our cospectrum results are consistent with the PDS results from other works (e.g. Grinberg *et al.* 2014).

Although we do not find any significant time lags between any of the energy bands, this is mostly due to the fact that our observation is not long enough. We estimate that $\sim 40\text{ ks}$ should be sufficient to detect time lags similar to those observed in the past in Cyg X-1, in this spectral state (e.g. Pottschmidt *et al.* 2000).

6.2. *GX 339-4*

GX 339-4 is another well known BHB. Discovered in 1973 by *OSO-7* (Markert *et al.* 1973), it has been observed frequently since then. It is a transient system, with long outbursts and periods of quiescence lasting up to several years (Belloni *et al.* 2005, and references therein). Monitored during its outbursts by *RXTE*, it showed one of the best examples of q-track pattern in the hardness-intensity diagram (Belloni *et al.* 2005; for the q-track pattern see Maccarone & Coppi 2003 and Dunn *et al.* 2010 for a comparison with other sources). Like many accreting BHs, its low/hard state PDS is dominated by two or more Lorentzian components, whose frequencies are positively correlated to count rate. As the source reaches the intermediate and soft state, the broad

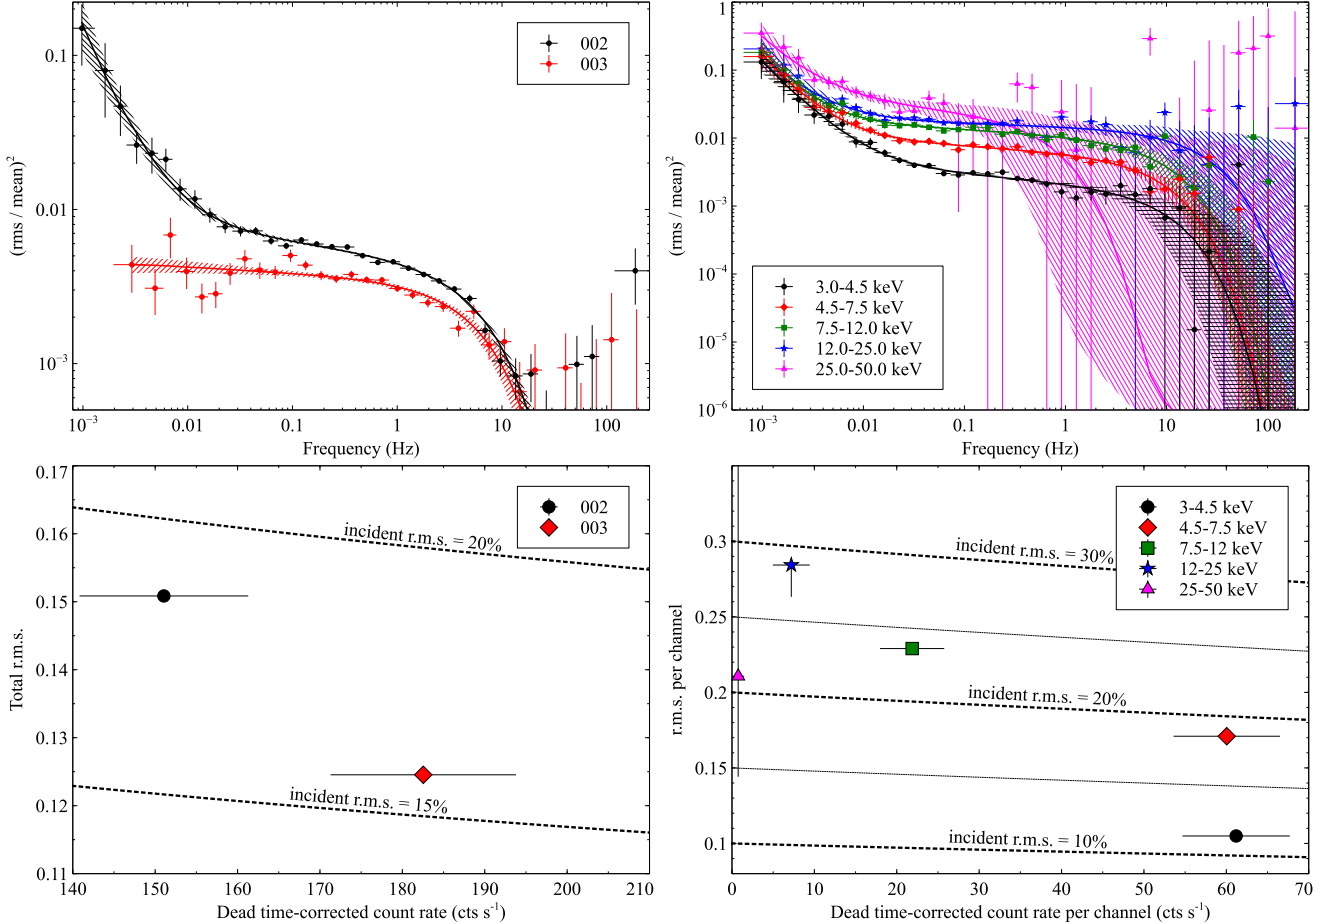


Figure 6. Cyg X-1: (Top Left) Comparison of the cospectrum in the two ObsIDs; the first observation is fitted with a `powerlaw+cutoffpl` model to account for the low-frequency excess, while the second with a simple `cutoffpl` model (the power law was not required by the fit). (Top Right) Cospectra in different energy ranges. There is a clear change of the cospectrum with energy, with a monotonic increase of r.m.s.. The cutoff frequency is consistent with being stable. Hatched regions represent the 90% errors ($\Delta\chi^2 = 2.706$) on the best-fit `powerlaw+cutoffpl` model for each energy range. (Bottom Left) Change of total r.m.s. in Cyg X-1 between the two observations. The points give the measured r.m.s., and the dashed lines show a grid of dead time-corrected r.m.s. values, following the methods shown in Figure 4. (Bottom Right) Change of r.m.s. at different energies in obs. 002. The r.m.s. error bars are smaller than the marker size in most cases.

Lorentzians move further to higher frequencies and various kinds of QPOs appear (see, e.g., Belloni et al. 2005; Casella et al. 2005), including high frequency QPOs. Thanks to this wealth of data, it has often been used, like Cyg X-1, to test new timing approaches. In its very high state, the signal at low energy ranges is generally coherent, while coherence is lost between low and high energies (Vaughan & Nowak 1997). Nowak et al. (1999b) measured coherence and lags in the hard state of the source, showing that the two Lorentzians are likely produced by independent, but internally coherent, processes. Cassatella et al. (2012) showed that hard time lags in this source are more likely to be produced by propagation of modulations of the accretion flow in the disk rather than from reflection.

NuSTAR observed GX 339–4 four times during the rise of the last outburst. During this outburst, the source remained in a canonical low/hard state without reaching the threshold luminosity that would have marked the transition to the intermediate and soft states. The spectral analysis will be published in a separate publication (Fürst et al 2014, *in prep.*), but in this Paper we will concentrate on the timing properties of the source, analyzed

with the methods described above.

The whole band and energy-dependent cospectra in the four different observations are shown in Figure 7. The approximate flux of the source increased almost linearly from ~ 10 to ~ 40 cts s⁻¹ in each detector between the first and the last observation. The general shape and the variation of the cospectrum is similar to what has been observed in the past for this source in the hard state (e.g. Nowak et al. 1999b; Belloni et al. 2005) and for other BHBs (McClintock & Remillard 2006), with two zero-centered Lorentzian components dominating the shape (L_1 and L_2 in Belloni et al. 2005), plus occasionally a third Lorentzian needed to account for an excess between them.

The cospectrum evolves with L_1 consistently increasing its characteristic frequency in subsequent observations (as luminosity increases), while L_2 does not change as much. This is consistent with what was previously reported by Belloni et al. (2005) using the PDS. The energy dependence of the cospectrum is very weak in the first observation, and tends to increase in the subsequent observations, while the total r.m.s. consistently decreases as the count rate increases. In the second, third and

fourth observation the high energy cospectrum has consistently a lower r.m.s. than the low energy one. Again, these cospectra results are consistent with previous PDS results (e.g. Nowak et al. 1999b).

From Figure 4, and the discussion about the much brighter Cyg X-1, it is clear that at these frequencies and for count rates changing from 10 to 40 ctss⁻¹ we do not expect a such a strong change in r.m.s., which must therefore be intrinsic to the source.

A more detailed discussion will be presented in Fürst et al., *in prep.*

7. CONCLUSIONS

This paper presents a method to exploit *NuSTAR*’s two independent detectors to work around the spurious effects resulting from dead time, enabling standard aperiodic timing analysis of X-ray binaries to be applied. As extensively studied in literature (e.g. van der Klis 1989; Vikhlinin et al. 1994; Zhang et al. 1995), dead time produces a spurious correlation between event times that strongly modifies the shape of the PDS. For frequencies $\nu \ll 1/\tau_d$, the only visible effect in PDSs is a slight decrease of the white noise level. In timing-oriented X-ray missions, dead time is very short (*RXTE*/PCA: 10 μ s, Jahoda et al. 2006; *ASTROSAT*/LAXPC: 10–35 μ s, Paul & Team 2009; *LOFT*: < 1 ns, Feroci et al. 2012) and the “interesting” range of frequencies for accretion (the highest being the dynamical timescales around neutron stars $\simeq 2000$ Hz) is well below $1/\tau_d$. In *NuSTAR*, τ_d is around 2.5 ms, right in the middle of the frequency range commonly analyzed in BH and NS binaries. Moreover, in order to obtain better precision in the measure of photon energies, the dead time is variable and depends on the event grades. Due to this, standard techniques for the modeling of the PDS shape (e.g. Vikhlinin et al. 1994; Zhang et al. 1995) are not applicable.

But *NuSTAR* has two independent focal plane modules that are read out by separate processors, and the dead time is therefore independent. If, instead of taking the PDS of each detector, we take the cospectrum (see Section 3) of the signals in the two detectors, one is able to cancel out most correlations produced by dead time and obtain a close proxy of a white noise-subtracted PDS.

We described in Sections 3 and 4 how to calculate the cospectrum and estimate the remaining effects of dead time on the signal, namely the modulation of the variance of data bins, that cannot be described using the standard Leahy et al. (1983) prescriptions, and the equivalent change of the observed r.m.s. at different frequencies and for different count rates. This estimate requires Monte Carlo simulations tailored to the source that is being analyzed, as described in Section 4.

We also applied this technique to the analysis of *NuSTAR* data of two BHBs, Cyg X-1 and GX 339–4. Cyg X-1 was observed in the soft state. We report a strong change of r.m.s. with energy, with a significant rise of the variability at higher energies. These cospectrum results are consistent with earlier PDS results, and extend the energy coverage with respect to earlier results from *RXTE* (Grinberg et al. 2014, e.g.). Comparing these results with the spectral modeling presented by Tomsick et al. (2014), it is apparent that the low energies are dominated by the disk emission, while at higher energies it’s the reprocessed component (Comptonized and

reflected) that dominates. This is consistent with what is found in many BHBs, and in particular to what has already been reported in this source (the “stable disk and unstable corona” picture, e.g. Churazov et al. 2001).

In GX 339–4, we observed the rising part of the outburst, in a standard low/hard state. The cospectrum shows that, as the outburst goes on, the r.m.s. generally decreases and the characteristic frequency of the first Lorentzian component used to described the cospectrum consistently rises. Also, this is consistent with what has been observed in the low/hard state of this and other sources using the PDS (Nowak et al. 1999b; Belloni et al. 2005).

We do not find evidence for time lags in the sources analyzed in this paper, due to the short observing time available, but the cross power density spectrum can be used to calculate time lags without major impact from dead time, if data are adequately filtered to avoid intervals with high background activity and in particular at the start and end of each good time interval.

The techniques presented in this paper can be used to perform most of the standard timing and spectral timing analysis used in literature, if the change of r.m.s. with frequency and count rate is properly taken into account (Section 4). For slow timing ($\nu \ll 1$ Hz), the dead time correction included in the official pipeline is still recommended, because it effectively corrects for all sources of dead time, including shielded events. For low count rates ($\lesssim 1$ ctss⁻¹) the standard PDS of the sum of the two lightcurves (what we called the “total PDS” above) is still advantageous, because the effects of dead time are negligible and the sensitivity of the total PDS, expressed as significance of detection of a standard variability feature, is $\sim \sqrt{2}$ higher than the cospectrum.

As an additional note, this technique can be used in any satellite that has completely independent detectors, and of course by using multiple satellites. *RXTE*, with its independent PCUs, would be a good candidate if not for the presence of so-called Very Large Events (VLE), particle events depositing more than 100 keV and occurring many times per second (e.g. Jones et al. 2008). These events often involve more than one PCU, invalidating the assumption of independence. The Large Area Xenon Proportional Counter (LAXPC) onboard the upcoming *ASTROSAT* mission (Agrawal 2006) will provide three independent detectors. Provided that the fraction of VLE-equivalent events is negligible, this technique will be applicable to LAXPC data.

MB and DB wish to acknowledge the support from the Centre National d’Études Spatiales (CNES) and the Centre National de la Recherche Scientifique (CNRS). CS acknowledges funding by the German BMWi under DLR grant numbers 50 QR 0801, 50 QR 0903, and 50 OO 1111. P.G. thanks STFC for support (grant reference ST/J003697/1). This work was supported under NASA Contract No. NNG08FD60C, and made use of data from the *NuSTAR* mission, a project led by the California Institute of Technology, managed by the Jet Propulsion Laboratory, and funded by the National Aeronautics and Space Administration. We thank the *NuSTAR* Operations, Software and Calibration teams for support with the execution and analysis of these




Figure 7. GX 339–4: Cospectra for the four pointings, in the 3–50 energy band (Left) and in two energy bands (Right). Best-fit models with 3 zero-centered Lorentzian components are also plotted. In the second plot, only the best fit models are plotted for clarity. There is a clear evolution of the cospectrum: first of all, the r.m.s. generally decreases as the outburst goes on; the Lorentzians, in particular below 1 Hz, generally increase their widths; also, while the two energy bands have the same cospectrum in the first observation, they separate more and more clearly in the following ones. We verified that this change is too strong to be produced by dead time alone (it can account for at most a 10% relative r.m.s. drop). In the latest two observations a hint of a power law component below 0.01 Hz also appears. It is not significant in our fits.

observations. This research has made use of the *NuSTAR* Data Analysis Software (NuSTARDAS) jointly developed by the ASI Science Data Center (ASDC, Italy) and the California Institute of Technology (USA). The timing analysis was executed with Matteo Bachetti’s Libraries and Tools in Python for *NuSTAR* Timing (MaLT PyNT). This code is available upon request. Most of the plots were produced with the Veusz software by Jeremy Sanders. The authors wish to thank Chris Done, Denis Leahy and Tomaso Belloni for very insightful discussions.

REFERENCES

- Agrawal, P. C. 2006, *ASR*, 38, 2989
 Arnaud, K. A. 1996, *ADASS*, 101, 17
 Artigue, R., Barret, D., Lamb, F. K., Lo, K. H., & Miller, M. C. 2013, *arXiv*, 1303.0248, 1303.0248
 Axelsson, M., Borgonovo, L., & Larsson, S. 2005, *A&A*, 438, 999
 Bachetti, M. et al. 2013, *ApJ*, 778, 163
 Barret, D., & Vaughan, S. 2012, *ApJ*, 746, 131
 Belloni, T., & Hasinger, G. 1990, *A&A*, 230, 103
 Belloni, T., Homan, J., Casella, P., van der Klis, M., Nespoli, E., Lewin, W. H. G., Miller, J. M., & Mendez, M. 2005, *A&A*, 440, 207
 Bendat, J. S., & Piersol, A. G. 2011, *Random Data: Analysis and Measurement Procedures* (Wiley)
 Böck, M. et al. 2011, *A&A*, 533, 8
 Boutelier, M. 2009, PhD thesis, PhD Thesis, CERN
 Boutelier, M., Barret, D., & Miller, M. C. 2009, *MNRAS*, 399, 1901
 Bowyer, S., Byram, E. T., Chubb, T. A., & Friedman, H. 1965, *Science*, 147, 394
 Cabanac, C., Roques, J.-P., & Jourdain, E. 2011, *ApJ*, 739, 58
 Casella, P., Belloni, T., & Stella, L. 2005, *ApJ*, 629, 403
 Cassatella, P., Uttley, P., Wilms, J., & Poutanen, J. 2012, *MNRAS*, 422, 2407
 Castro-Tirado, A. J., Brandt, S., & Lund, N. 1992, *IAU Circ.*, 5590, 2
 Churazov, E., Gilfanov, M., & Revnivtsev, M. 2001, *MNRAS*, 321, 759
 Cui, W., Zhang, S. N., Jahoda, K., Focke, W., Swank, J. H., Heindl, W. A., & Rothschild, R. E. 1997, *The Transparent Universe*, 382, 209
 Davies, R. B., & Harte, D. S. 1987, *Biometrika*, 74, 95
 Dunn, R. J. H., Fender, R. P., Körding, E. G., Belloni, T., & Cabanac, C. 2010, *MNRAS*, 403, 61
 Fabian, A. C. et al. 2009, *Nat.*, 459, 540
 Feroci, M. et al. 2012, *Experimental Astronomy*, 34, 415
 Foster, R. S., Waltman, E. B., Tavani, M., Harmon, B. A., Zhang, S. N., Paciesas, W. S., & Ghigo, F. D. 1996, *Astrophysical Journal Letters* v.467, 467, L81
 Fürst, F. et al. 2013, *ApJ*, 779, 69
 —. 2014, *ApJ*, 780, 133
 Gandhi, P. et al. 2010, *MNRAS*, 407, 2166
 Gentle, J. E. 2003, *Random Number Generation and Monte Carlo Methods* (Springer)
 Gilfanov, M., Churazov, E., & Revnivtsev, M. 2000, *MNRAS*, 316, 923
 Gleissner, T., Wilms, J., Pottschmidt, K., Uttley, P., Nowak, M. A., & Staubert, R. 2004, *A&A*, 414, 1091
 Grinberg, V. et al. 2014, *A&A*, 565, 1
 Harrison, F. A. et al. 2013, *ApJ*, 770, 103
 Jahoda, K., Markwardt, C. B., Radeva, Y., Rots, A. H., Stark, M. J., Swank, J. H., Strohmayer, T. E., & Zhang, W. 2006, *ApJ Supp. Ser.*, 163, 401
 Jin, C., Done, C., Middleton, M., & Ward, M. 2013, *MNRAS*, 436, 3173
 Jones, T. A., Levine, A. M., Morgan, E. H., & Rappaport, S. 2008, *ApJ*, 677, 1241
 Kaspi, V. M. et al. 2014, *ApJ*, 786, 84
 Körding, E., & Falcke, H. 2004, *A&A*, 414, 795
 Leahy, D. A., Darbro, W., Elsner, R. F., Weisskopf, M. C., Kahn, S., Sutherland, P. G., & Grindlay, J. E. 1983, *ApJ*, 266, 160
 Lewin, W. H. G., van Paradijs, J., & van der Klis, M. 1988, *SSRv*, 46, 273
 Maccarone, T. J., & Coppi, P. S. 2003, *MNRAS*, 338, 189
 Marinucci, A. et al. 2014, *MNRAS*, 440, 2347
 Markert, T. H., Canizares, C. R., Clark, G. W., Lewin, W. H. G., Schnopper, H. W., & Sprott, G. F. 1973, *ApJ*, 184, L67
 McClintock, J. E., & Remillard, R. A. 2006, in *Compact stellar X-ray sources*, ed. J. E. McClintock & R. A. Remillard, 157–213
 Miller, J. M. et al. 2013, *ApJL*, 775, L45
 Miyamoto, S., Kimura, K., Kitamoto, S., Dotani, T., & Ebisawa, K. 1991, *ApJ*, 383, 784
 Miyamoto, S., & Kitamoto, S. 1989, *Nat.*, 342, 773
 Miyasaka, H. et al. 2013, *ApJ*, 775, 65
 Morgan, E. H., Remillard, R. A., & Greiner, J. 1997, *ApJ*, 482, 993
 Mori, K. et al. 2013, *ApJL*, 770, L23
 Natalucci, L. et al. 2014, *ApJ*, 780, 63
 Nowak, M. A., & Vaughan, B. A. 1996, *MNRAS*, 280, 227
 Nowak, M. A., Vaughan, B. A., Wilms, J., Dove, J. B., & Begelman, M. C. 1999a, *ApJ*, 510, 874
 Nowak, M. A., Wilms, J., & Dove, J. B. 1999b, *ApJ*, 517, 355
 Pahari, M., Neilsen, J., Yadav, J. S., Misra, R., & Uttley, P. 2013, *ApJ*, 778, 136




Figure 8. GRS 1915+105: (Left) Cospectrum before (black) and after (red) clean up of the n th harmonics of the housekeeping operations (at 1 and 0.125 Hz)⁷. (Right) evolution of the 1.5 Hz QPO (and a faint first harmonic) in the dynamical cospectrum during the observations. Each time segment represents 64s of gap-free data. Colors and sizes of the points are both a proxy of power, in arbitrary units.

Papadakis, I. E., Nandra, K., & Kazanas, D. 2001, *ApJ*, 554, L133
 Paul, B., & Team, L. 2009, in *Astrophysics with All-Sky X-Ray Observations*, 362
 Pottschmidt, K., Wilms, J., Nowak, M. A., Heindl, W. A., Smith, D. M., & Staubert, R. 2000, *A&A*, 357, L17
 Reig, P., Belloni, T., van der Klis, M., Mendez, M., Kylafis, N. D., & Ford, E. C. 2000, *ApJ*, 541, 883
 Remillard, R. A., Munro, M. P., McClintock, J. E., & Orosz, J. A. 2003, *AAS*, 7, 648
 Revnivtsev, M., Gilfanov, M., & Churazov, E. 1999, *A&A*, 347, L23
 —. 2000, *A&A*, 363, 1013
 Risaliti, G. et al. 2013, *Nat.*, 494, 449
 Strohmayer, T. E. 2001, *ApJ*, 554, L169
 Timmer, J., & Koenig, M. 1995, *A&A*, 300, 707
 Tomsick, J. A., & Kaaret, P. 2001, *ApJ*, 548, 401
 Tomsick, J. A. et al. 2014, *ApJ*, 780, 78
 Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, *arXiv*, 6575, 1405.6575

Uttley, P., Wilkinson, T., Cassatella, P., Wilms, J., Pottschmidt, K., Hanke, M., & Böck, M. 2011, *MNRAS Let.*, 414, L60
 van der Klis, M. 1989, in *Timing Neutron Stars: proceedings of the NATO Advanced Study Institute on Timing Neutron Stars held April 4-15*, 27
 Vaughan, B. A., & Nowak, M. A. 1997, *ApJL*, 474, L43
 Vaughan, S. 2013, *arXiv*, 1309.6435v1
 Vikhlinin, A., Churazov, E., & Gilfanov, M. 1994, *A&A*, 287, 73
 Walton, D. J. et al. 2014, *arXiv*, 5620, 1404.5620
 Weisskopf, M. C. 1985, in *Talk presented at the Workshop Time Variability in X-ray and Gamma-Ray sources*, Taos, NM, USA
 Zhang, W., Jahoda, K., Swank, J. H., Morgan, E. H., & Giles, A. B. 1995, *ApJ*, 449, 930
 Zoghbi, A. et al. 2014, *ApJ*, 789, 56

APPENDIX

GRS 1915+105: TIMING ANALYSIS FROM INSTRUMENT COMMISSIONING

GRS 1915+105 is a very bright persistent microquasar, discovered in 1992 by *Granat* (Castro-Tirado et al. 1992) which shows accretion states that depart from the usual behavior of Galactic BHs, being associated with very high-Eddington fractions. The variability of this source has been extensively studied in the past, in particular for the presence of very strong low frequency (<0.1 Hz), medium-low frequency (0.5–10 Hz), medium-high frequency (~ 40 and 67 Hz) and high frequency (113 and 165 Hz) QPOs (e.g. Morgan et al. 1997; Strohmayer 2001; Remillard et al. 2003; Pahari et al. 2013). In our observation, the medium-low QPO is the only one observed to very high significance. Its behavior has been studied in great detail (e.g. Reig et al. 2000) and its properties are known up to very high energies thanks to *RXTE*/HEXTE observations (Tomsick & Kaaret 2001). This source is discussed as a reference for timing analysis on sources observed very early in the mission. The first version of the flight software, in fact, introduced a large number of spurious modulations in the signal that we will discuss and show a way to work around.

During the observatory commissioning phase prior to commencement of science observations, *NuSTAR* observed the galactic BHB GRS 1915+105 in its *plateau* state (Foster et al. 1996). The spectral analysis of that observation was presented by Miller et al. (2013). The version of the flight software that was installed at that time on *NuSTAR* executed housekeeping operations at regular intervals (1 or 8s), producing periodic dead time. Therefore, the raw power spectrum and the CPDS of this observation are plagued by spurious features at 1 Hz and 0.125 Hz as well as their harmonics and sub-harmonics (see Figure 8). In this Section we show a technique to clean up the cospectrum and perform a basic analysis of the strong QPO that this source.

To “clean” the cospectrum from these spurious features, we first produced a cospectrum on a time scale t_{fft} of 128 s. The maintenance operations were very regular, and so in this cospectrum the corresponding features were very sharp and limited to a single bin each. We replaced the value of each of these bins with a random value calculated from the mean of the two closest bins and their standard deviation, as shown in Figure 8. The possible spurious effects of this imperfect statistical treatment of the simulated bin value are negligible due to the small frequency bin (simulated data are only one bin every $1/t_{\text{fft}} = 1/128$ Hz) and the rebinning applied in the following steps of the analysis. The real and imaginary part were treated independently.

From this cleaned CPDS we calculated the cospectrum. An example is shown in red in Figure 8. The cospectrum shows a very strong and variable QPO at ~ 1.5 Hz, with r.m.s. = 10(1)% (from the best-fit Lorentzian amplitude, and corrected for the expected $\sim 20\%$ drop at ~ 100 cts s $^{-1}$, Equation 5). This feature has previously been observed in this spectral state (e.g. Pahari et al. 2013). In Figure 8 we show that with these techniques it is possible in principle to also apply techniques of QPO tracking (a dynamical cospectrum) given a sufficient r.m.s.. A hint of the harmonic of this QPO is also observable, both in the single and in the dynamical cospectra. No other significant QPOs are detected. This source is known to show high-frequency QPOs in *RXTE* observations (Morgan et al. 1997; Remillard et al. 2003). These features are not observed in our dataset. We calculate a rough upper limit of $\sim 5\%$ for additional QPOs between 50 and 100 Hz. This upper limit is higher than the observed r.m.s. of QPOs in this frequency range observed in the past (e.g. $\sim 1.5\%$ in the full *RXTE* band in Morgan et al. 1997), and our non-detection might be largely due to the low statistics of this dataset.